

Silicon Lifecycle Management: Actionable Silicon Insights Through Intelligent Measurement and Analysis

Executive Summary

Semiconductors have always been challenging to develop, with many waves of innovation in electronic design automation (EDA) tools and fabrication technologies barely keeping ahead of ever-growing design size and sophistication. Once again, the industry has reached a tipping point. The combination of increasing chip and system complexity, coupled with higher expectations for product performance and longevity, has defined new boundary limitations in designing, manufacturing and deploying semiconductors. Incremental improvements to traditional methods will not be enough to move forward; a new approach is required.

Silicon Lifecycle Management (SLM) is an emerging paradigm within the industry that will make product development and deployment more deterministic. SLM consists of two primary steps:

- Deploying monitors and structures embedded within silicon designs to gain insight on how devices are made and how they then perform in-field
- Gathering and analyzing data at every opportunity throughout the lifecycle of silicon devices to provide powerful analytics that enable in-design, in-test and in-field observations and device improvements to be made

SLM enables the semiconductor design community to exploit a “monitor, analyze and optimize” philosophy. The key benefits are improvements in quality, performance and reliability of silicon systems, enabling predictive maintenance and failure prediction. These lifecycle insights provide significant capability and performance benefits to forward-thinking users, including those in hyperscaling, consumer and automotive applications.

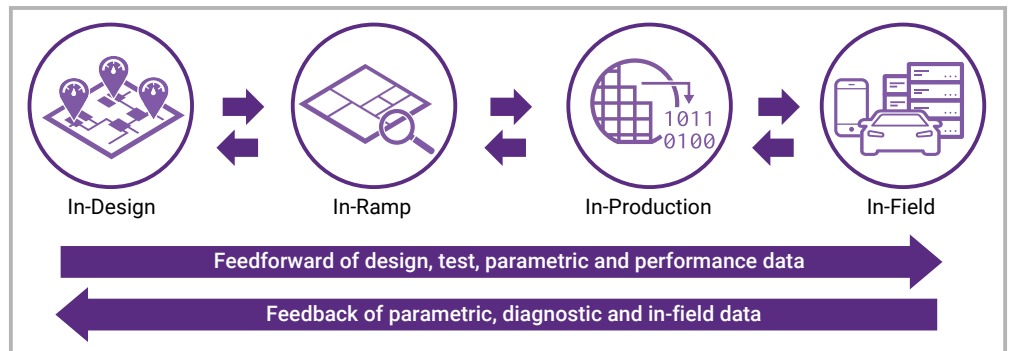


Figure 1: Silicon Lifecycle Management concept

Synopsys is developing new and insightful SLM monitoring and analytics technologies, under a single platform, that resonate with customer and partner ecosystems. This is one of the most exciting growth sectors of the semiconductor industry and Synopsys has emerged as a significant contributor in this space with its Silicon Lifecycle Management Family.

Background

The semiconductor industry is facing new challenges posed by the accelerated scaling of device and system complexity. For each new technology node, transistor densities continue to increase. This provides a great opportunity for adding technology capabilities, but it also great challenges. With increasing transistor density comes more significant variability in the manufacturing process. This broadens the design envelope unless process variability across the die can be measured using monitoring structures.

Increased design density also produces greater current density and, therefore, power density. This presents challenges from diminished advanced node voltage supply levels varying due to dynamic workload activities. The heat dissipated throughout the die due to increased transistor densities is also a challenge, creating hotspots that need to be managed to tame the power conditions and improve the chip’s reliability. Other forms of increasing complexity include:

- Chip complexity—placement and routing complexity
- System complexity—die stacked in multi-chip module (MCM) or 2.5D/3D arrangements
- Hardware and software design complexity and interactions—the way the SoC design reacts to stimulus, executes instructions and transacts data is not entirely deterministic through the device’s lifetime as firmware and software are upgraded

Greater density and complexity increase the likelihood of physical failure. This, combined with expectations for zero faults in manufacturing and increased longevity in-field, presents additional challenges. Circuit testing must extend beyond the manufacture, bring-up and production test phases to normal operation phases in mission mode in-field. Continuous testing throughout the product’s lifecycle is required.

Better and more efficient designs are essential. Figure 2 represents the silicon and software design costs for an advanced performance multicore system-on-chip (SoC) silicon solution. Not every design will reach these cost levels. If the design is a first-time effort, is from one of the first companies to design at a new process node and maximizes all the design parameters (highest possible gate count, fastest speed and largest die area), costs like those shown in the graph can be incurred. Designs with device parameters that are not pushed to the limit may incur lower design costs.




Figure 2: Rising silicon and software design costs

Along with the rising costs associated with each stage of a chip’s lifecycle, there are growing amounts of data generated. These can include design data, manufacturing parametric data, test data and in-field data. Data coherency across each of these stages, different users and multiple vendors is an emerging area for the industry. Work in this field may lead to alliances and standards being defined for data capture, formatting, contextualization, consolidation, compression, telemetry, storage and security. These needs, along with the lifecycle analytics that leverage the data, drive the important emerging SLM domain.

New Paradigm for System SLM

Increasing system complexity, compounded with the high performance and reliability expectations of data center, automotive and consumer applications, places high demands on chip developers. In response, the semiconductor industry is required to introspectively examine each process step in the creation and deployment of silicon devices. Such examination will deliver excellent value with higher degrees of data visibility and analytics that lead to optimization actions:

- In-silicon visibility: By embedding monitors, instruments or agents into the design, it is possible to act on conditions of how a device has been made and how it responds to stimuli. There is a strong desire for measurement data that is increasingly more accurate, more distributed and more responsive. As well as the application of embedded structures to measure parametric silicon information, new monitor types are being developed to uncover further information, becoming additional sources for the analysis of static and dynamic data.
- In-test visibility: This includes the application of analytics during test phases to increase the visibility of failing devices and change what is acceptable as a known-good-die, providing comprehensive screening of silicon and expanding the criteria for what is accepted.
- In-field visibility: Deployment of silicon into systems in the field, large sample sets, analytics across entire product ranges or fleets and visibility of longer-term trends are all highly desirable.
- Data accessibility: Users want insightful and meaningful data that has been compressed and contextualized, with web-based graphical user interfaces (GUIs) showing this data appropriately across an organization, from engineers to executives. This ensures that users do not “drown” in raw, voluminous data.

SLM provides all these aspects of data visibility and performs the required analysis on this data. Thus, lifecycle management for semiconductor-based systems is emerging as one of the key areas of investment for the chip industry. Its concepts must be ingrained within the language of design, development, manufacture, production test and quality teams. Successful silicon lifecycle management integrated systems must be built on a unified database solution, capturing and storing pertinent information from each lifecycle stage, as shown in Figure 3. An SLM vendor must provide both native and third-party database solutions depending upon the user’s attitude toward data sensitivity.

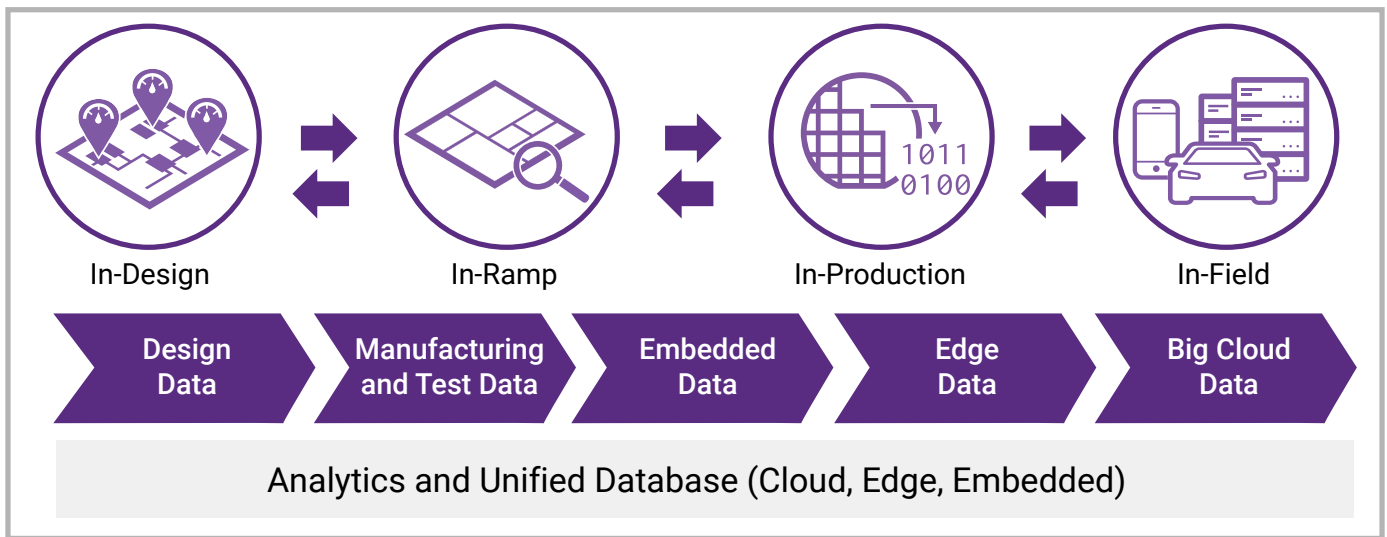


Figure 3: Integrated SLM flow across lifecycle stages

The SLM system must facilitate an end-to-end solution, driven by industry leaders coupled with partner innovation. This must include:

- In-chip monitoring of environmental, structural and functional conditions
- Automated monitor selection and insertion into the design
- Accelerated product ramp by identifying systematic post-silicon design issues
- Automated analysis and production control of high-volume manufacturing test data across the entire supply chain
- Adaptive high bandwidth testing over functional interfaces

- Embedded edge analytics solutions within devices operational in the field
- An ecosystem to enable data telemetry, manipulation, storage, security and analysis across monitors

The SLM system must also be flexible, able to grow and scale over time. It must allow easy adoption of new structures and monitors. It must monitor various aspects of silicon parameters, including those that change over time, such as aging and degradation. The system should accommodate novel “instruments” or “agents” developed by the SLM vendor or third-party partners.

Key Benefits of SLM

By improving lifecycle visibility, enriched analytics and reporting, SLM enables the semiconductor design community to explore enhanced versions of established control mechanisms. This optimizes device operation for either power or data throughput performance, under the “monitor, analyze and optimize” philosophy mentioned earlier. Key benefits include improvements in the quality, performance and reliability of silicon systems.

Lifecycle management is not achieved by deploying one single product or tool but instead through the integration of constituent components, the combination of which will depend upon user product requirements. High volume consumer applications may require a solution focused on reducing design constraint pessimism by gathering parametric feedback from silicon. For example, wafer acceptance testing (WAT) results can help predict chip binning at wafer sort and final test. Parametric test results can show where the tail of the process distribution fails the test limits and where there are statistical outliers that fall within limits.

Hyperscaling applications may require highly granular thermal management within processor cores or clusters throughout the die to optimize power performance. Even small (micro) power savings can scale to macro savings in a massive cloud server configuration. For example, a modest enhancement to thermal sensing accuracy reducing power by less than one penny per processor chip per hour can save millions of dollars per year in a large data center. Now that servers cost more to run over their lifetime than their initial purchase price, any reduction in power consumption is significant.

Automotive applications for semiconductors present their own unique set of challenges. Through continuous assessment of aging and degradation factors, such as heat stress, it may be possible to predict when in-vehicle electronic systems require maintenance or replacement. If automotive systems are more predictive for fatigue and failure, there are opportunities to incorporate commercial-grade chips with the understanding that they will statistically fail earlier but in a more predictable way. This may allow increased use of lower-cost solutions with limited design-for-test (DFT) and design-for-manufacturing (DFM) elements rather than expensive hyper-reliable systems. Figure 4 provides a visualization of this effect.




Figure 4: Advantages of SLM for high reliability automotive applications

Synopsys SLM Family

SLM is an increasingly important element of the production, manufacturing and deployment flow for SoCs and other advanced semiconductors. Synopsys has emerged as a key player in SLM with its Silicon Lifecycle Management (SLM) Family, a solution that meets all the requirements and provides all the benefits discussed thus far, including a Unified Database. In addition, Synopsys Silicon Lifecycle Management Family has unique capabilities that position Synopsys as an industry thought leader. Figure 5 shows the key components of this end-to-end solution. These include:

- Embedded monitoring IP, including PVT and path margin
- Digital Design Platform, silicon-aware design flow improvement
- Yield Explorer (production ramp, bring-up and diagnostics)
- SiliconDash (high volume product analytics screening and production control)
- SLM HS Access IP and TestMAX™ ALE (adaptive high-bandwidth testing)
- SLM Optimizer (performance tuning)

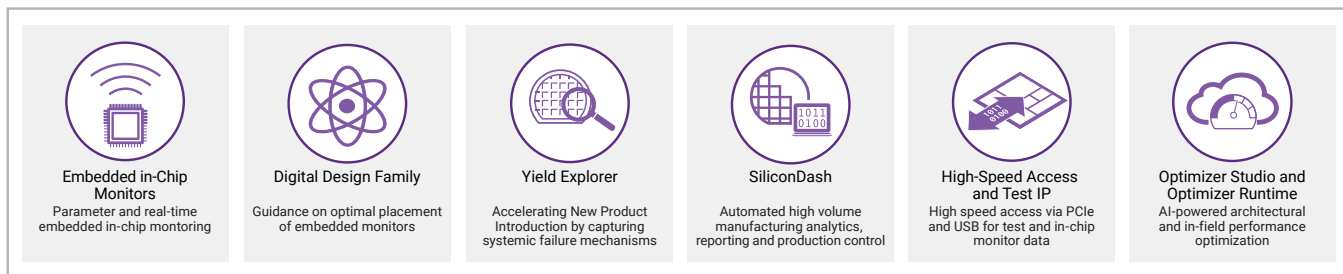


Figure 5: Synopsys Silicon Lifecycle Management Family

The SLM process starts during the architectural phase with the selection of a high-level hardware design that optimizes system performance. SLM Optimizer uses artificial intelligence (AI) techniques to explore the design space and model various design options quickly and automatically. Compute performance is usually measured by standard benchmark programs, and Optimizer can co-optimize both software compiler settings and hardware to minimize benchmark runtime. Architects can also run sensitivity analysis to explore the complex tradeoffs between hardware, software and performance.

Designers then have the opportunity to embed monitor IP within the chip architecture and the detailed design, allowing dynamic thermal and supply voltage data to be gathered for in-test data generation and real-time analysis for optimization scheme control loops. The embedded IP can measure the extent of silicon process variability and assess the relationship between safe timing margins for digital logic and its voltage supply levels. The measured data can be used to tighten design margins and enable real-time supply level optimization. Combined with accurate analytics, such approaches minimize pessimism and over margining in the design process.

There is also the opportunity for parametric data to be harvested across high volumes of production chips, allowing simulations to use updated highly representative timing models and the machine-learning selection of Monte Carlo simulations. This reduces simulation time and improves the power, performance and area (PPA) of the design. Looking ahead, there is the opportunity for further monitor innovation, looking at various stress mechanisms, droops on supplies, critical path margins and localized power consumptions within the chip. Design information can be added to the Unified Database, allowing design-based assessment for each lifecycle stage.

By having a dedicated analytics engine resource, chip-level optimization schemes and analytics can be developed to provide low-latency decision-making and enable responsive power management. The Synopsys Embedded Learning Engine (ELE) allows voluminous monitor data to be interpreted, reducing the amount being passed up the software stack to the industry's most insightful, information-rich analytics. It is vital to pass only the meaningful data rather than large quantities of data that consume high bandwidth and overhead in telemetry. The ELE combined with software allows monitor data to be formatted in convenient forms for analytics.

The Synopsys Digital Design Family solution offers a way to feed the parametric assessment of silicon produced back into the design process. By evaluating silicon at high volume levels, an understanding of where the manufacturing process is centered (such as

whether slow, typical or fast silicon speeds dominate) can be formed. This data can be used to tune and calibrate models used in simulation. Thus, the design process is tightened and made less pessimistic, creating an opportunity for silicon area (through better-tuned logic synthesis) and power reduction.

By inserting DFT test structures for scan and built-in self-test (BIST) as a part of the design flow, significant value is provided during chip production to screen for “good” die (as measured by functionality, performance and reliability). Yield Explorer provides design-centric analysis and reporting for DFT and failure analysis for production when ramping to volume. SiliconDash provides production-centric analysis and control to improve quality and productivity, handles large data sets and provides reports via an efficient online user interface. Both Yield Explorer and SiliconDash connect to the SLM Unified Database, offering a coordinated data repository for interrogation by test, quality and product engineers.

SLM also brings value to the system bring-up process. As components are assembled within single or multi-die packages, placed on printed circuit board (PCB) substrates or encompassed within a final product, the traceability of silicon to ensure compatibility across the connected components creates a demand for in-line testing aware of the test performance of individual die. Possessing this information allows the product to be confidently assembled by applying tests that span from individual die to complex assembled systems. The Synopsys SiliconDash solution provides an analytics capability that communicates faults when they occur via online reporting to production teams. This is also supported by Synopsys TestMAX, allowing for in-production and in-field tests to be conducted, flagging issues at each lifecycle stage.

Software embedded into the target system in-field supports local analysis and communication to the SLM Unified Database. An in-field test is possible by embedding the Synopsys High-Speed Access & Test (HSAT) IP interface solution. Using in-situ high-bandwidth standard interfaces such as PCIe and USB, test data can be efficiently presented to the device and resulting data can be shifted off the chip. This makes continual lifetime testing viable, allowing degradation or aging of devices to be monitored in-field over fleets of devices. Such schemes include logic and memory test, allowing in-field on-chip memory diagnosis (OCMD) to be performed periodically during mission mode or during each power-up sequence of the device.

Software running in the target system includes SLM Optimizer Runtime, the autonomous real time performance tuning solution. It accelerates performance by applying low-level accelerations and tailoring the many system and application settings to work in concert with the currently running workloads. An AI-powered dynamic tuning agent runs in the background, continuously monitoring system metrics and optimizing performance. As shown in Figure 6, there are three dimensions to its analysis and tuning. This approach can be applied to any software stack through the use of plugins that define the settings available for tuning and the metrics used to measure system performance. All programs and applications run using the existing binaries, with no code changes or recompilations required.

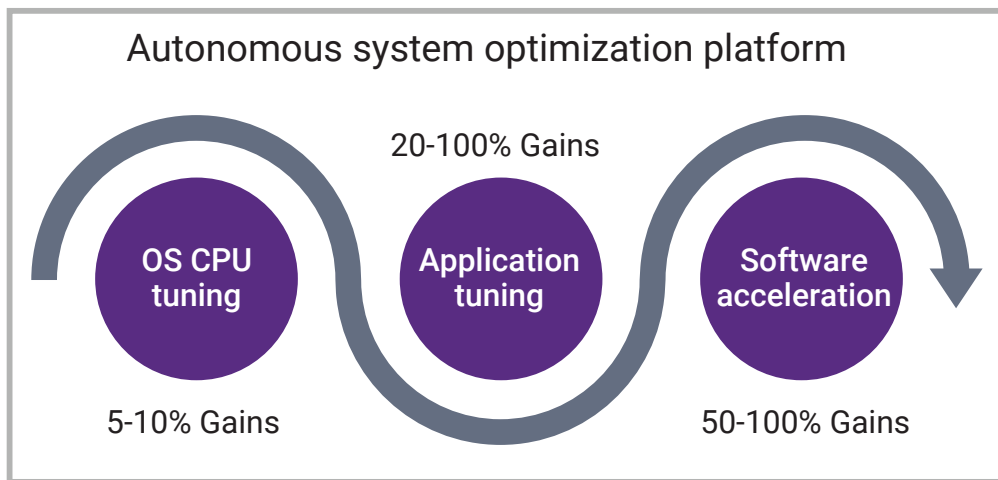


Figure 6: Dimensions of SLM Optimizer tuning

Synopsys Silicon Lifecycle Management Family is highly flexible. Able to grow and scale, it will encourage the adoption of new structures and monitors, looking at various aspects of silicon parameters (including those ranging over time, such as aging and degradation). Synopsys SLM will be able to accommodate novel “instruments” or “agents” developed by third-party partners or by Synopsys.

SLM Use Cases

SLM is a powerful and flexible concept, and Synopsys SLM leverages this flexibility for application in many different scenarios or use cases. One such use case is Design Links, in which parametric and critical path margin information from silicon manufacturing is fed back into the design flow via Digital Design Family. This silicon-to-design calibration provides several valuable benefits:

- Reduced design margin, reducing simulation over-pessimism
- Optimization of design libraries
- Optimization of ramp to volume by discovering precisely the cells contributing to systematic device failure
- In-field silicon failure debug, assisted by comparing failing transition delay paths within silicon against design tool timing reports

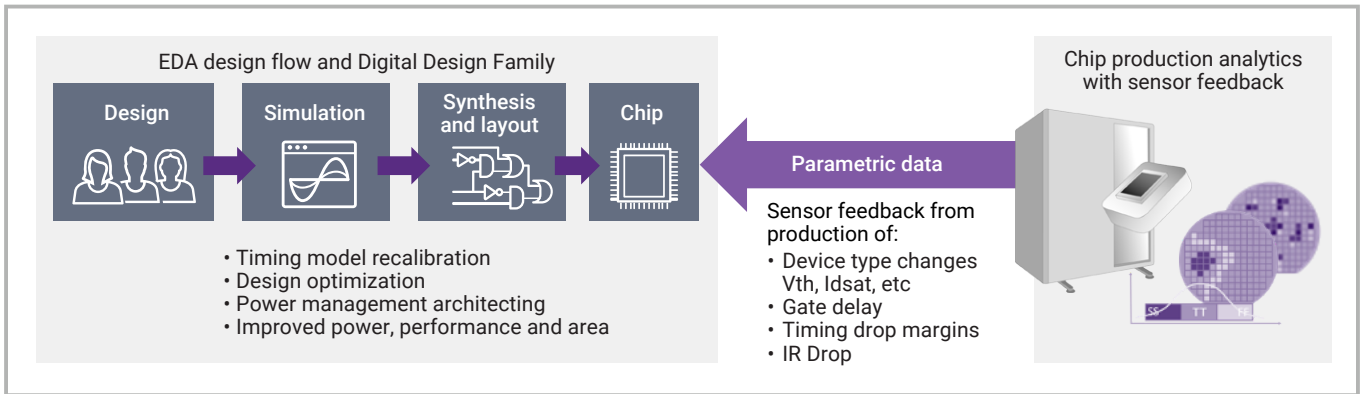


Figure 7: Synopsys Digital Design Family establishes link for silicon to design calibration

Another representative use case is process, voltage and temperature (PVT) aware testing, a re-purposing of the existing mission-mode monitors for chip production test on automatic test equipment (ATE). This provides two main opportunities. First, by measuring thermal and voltage supply conditions during test (scan and BIST), a richer diagnosis of device failure can be attained. A device failure may be attributable to anomalies such as increased die temperature conditions (localized hotspots) or excessively low supply levels (IR drop caused by either high chip dynamic activity or the intrinsic nature of the design).

In addition, die considered "good" may exhibit temperature and supply voltage anomalies as reported by the monitors. In this case, decisions can be made to screen such devices using more aggressive ATE selection schemes. This screening may be more suitable for some end applications, such as high reliability automotive.

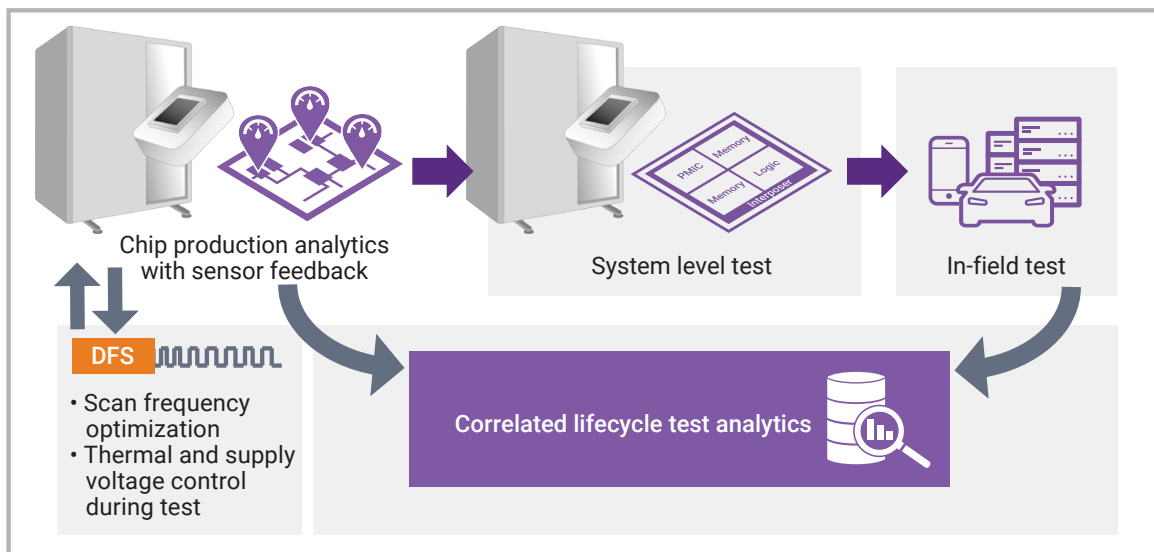


Figure 8: Monitor aware testing

Finally, SLM supports in-field predictive maintenance and failure. By amalgamating the far-reaching array of analysis components within the SLM platform, signatures and trends about chip behavior can be assessed to build predictive models of silicon reliability. Chip condition information extracted during in-field operation creates a picture of usage, the conditions internal to the device and the external conditions to which it is subjected. Examples include the long-term assessment of thermal and voltage stress applied to the silicon.

Summary

Semiconductor design, manufacture and deployment become more challenging and less deterministic with advanced nodes and demanding applications. SLM addresses these concerns by embedding monitors and other structures within the silicon to monitor device behavior all the way to in-field usage, and then collecting and analyzing the data. Synopsys is leading the industry with the SLM platform, a unified integration of multiple proven tools and technologies. It collects monitor data in an intelligent and efficient fashion and provides the most powerful and valuable analytics available in the market.

Chip and system developers and end product users leveraging Synopsys SLM will reap huge benefits over their competitors, including more efficient design, better manufacturing productivity and predictive reliability. Predictions of in-field chip degradation or failure are extremely valuable for end users by enabling corrective action before sudden and catastrophic system failures. Given all these benefits, it is just a matter of time before adopting and deploying Synopsys SLM is standard procedure on every project.